

Bounds for rating override rates

Dirk Tasche*

First version: March 10, 2012

This version: July 13, 2012

Overrides of credit ratings are important correctives of ratings that are determined by statistical rating models. Financial institutions and banking regulators agree on this because on the one hand errors with ratings of corporates or banks can have fatal consequences for the lending institutions and on the other hand errors by statistical methods can be minimised but not completely avoided. Nonetheless, rating overrides can be misused in order to conceal the real riskiness of borrowers or even entire portfolios. That is why rating overrides usually are strictly governed and carefully recorded. It is not clear, however, which frequency of overrides is appropriate for a given rating model within a predefined time period. This paper argues that there is a natural error rate associated with a statistical rating model that may be used to inform assessment of whether or not an observed override rate is adequate. The natural error rate is closely related to the rating model's discriminatory power and can readily be calculated.

KEYWORDS: Credit rating, rating override, discriminatory power, accuracy ratio, misclassification rate.

1 Introduction

Overrides of credit ratings generated by statistical models are a somewhat controversial topic. Financial institutions and regulators alike acknowledge that statistical models while being useful for the acceleration of rating processes and quantification of rating results, in principle are still inferior to careful expert assessments of creditworthiness. Nonetheless, even experienced credit experts could undeliberately be biased to under- or overestimate the creditworthiness of certain borrowers or groups of borrowers.

The following comment by the Basel Committee on Banking Supervision ([BCBS, 2006](#), extract from paragraph 417) confirms and extends these observations:

“Credit scoring models and other mechanical rating procedures generally use only a subset of

*E-mail: dirk.tasche@gmx.net

The author currently works at the UK Financial Services Authority. The opinions expressed in this paper are those of the author and do not necessarily reflect views of the Financial Services Authority.

available information. Although mechanical rating procedures may sometimes avoid some of the idiosyncratic errors made by rating systems in which human judgement plays a large role, mechanical use of limited information also is a source of rating errors. Credit scoring models and other mechanical procedures are permissible as the primary or partial basis of rating assignments, and may play a role in the estimation of loss characteristics. Sufficient human judgement and human oversight is necessary to ensure that all relevant and material information, including that which is outside the scope of the model, is also taken into consideration, and that the model is used appropriately.”

While it is generally expected that overrides improve the quality of rating assignments¹, the Basel Committee has a clear view on the need to monitor overriding activity (BCBS, 2006, paragraph 428):

“For rating assignments based on expert judgement, banks must clearly articulate the situations in which bank officers may override the outputs of the rating process, including how and to what extent such overrides can be used and by whom. For model-based ratings, the bank must have guidelines and processes for monitoring cases where human judgement has overridden the models rating, variables were excluded or inputs were altered. These guidelines must include identifying personnel that are responsible for approving these overrides. Banks must identify overrides and separately track their performance.”

This paper is about analysing the performance of the overrides of the output of statistical rating models. How to do this is by no means obvious. Look for instance at the case where a borrower’s rating is suggested to be poor by a statistical model but is overridden to a high quality rating grade. Assume that the borrower afterwards defaults within the following twelve months. At first glance, this incident might be considered clear evidence of an unjustified override. It could, however, just be an occurrence of bad luck. Only if we observed a significant number of such outcomes of overrides we would be able to draw a conclusion on systematic bias of the overrides.

From this example, we can derive a first performance criterion for rating overrides: The discriminatory power of a rating system post overrides should not be lower than the discriminatory power of the system without the overrides. Another criterion, also based on a significant sample of overrides, relates to the tendency of the overrides: If the majority of the overrides is towards rating grades indicating poorer credit quality one has to investigate whether or not this is caused by underestimation of the probabilities of default (PDs)² associated with the grades of the rating system. Careful analysis should be applied to the reasons quoted as causing the overrides: They could give indications of risk drivers not captured by the statistical model or no longer being predictive of creditworthiness.

Financial institutions should also look at the frequency of overrides applied to the outcomes of rating models. Intuitively, observing a high frequency of overrides might indicate that something is wrong with the model. This is why most financial institutions monitor the override frequencies

¹The empirical evidence of the impact of overrides on rating performance is not unambiguous. Thus, in a recent study [Brown et al. \(2012\)](#) found that “overall, our results suggest that the widespread use of discretion by loan officers may not result in more accurate assessments of the creditworthiness of borrowers”. This finding, however, might be a consequence of the fact that most of the banks providing data for the study were not subject to the minimum requirements of the Basel II Internal Ratings Based Approach.

²In the following we will often use the acronym PD for probability of default.

of their wholesale rating systems and, typically, trigger corrective actions when pre-defined bounds for the override rates are exceeded.

What, however, does “high frequency” mean? Would it be 10% or rather 40% of the ratings assigned within a year? By intuition one would say that the “critical” frequency of overrides to indicate problems with a statistical rating model should be related to the discriminatory power of the model. As demonstrated in figure 1, a model with low power needs more corrections – in the shape of overrides – than a model with high power because the overlap between the rating profiles of defaulting and solvent borrowers is larger for low power rating models. In this paper, we suggest a method to determine bounds for override rates that should not be exceeded if all were right with the underlying statistical model. We then investigate the link between the discriminatory power of the statistical rating model and the proposed rating override bounds, finding that indeed the bounds are the tighter the more powerful the model is.


We also demonstrate how ex ante estimates of discriminatory power can be inferred from the PD curves associated with statistical rating models. Such estimates may then be used to derive the override rate bounds needed for effective monitoring. This is particularly useful in the case of low default portfolios where – due to the lack of default observations – monitoring and validation to a large extent must be based on the analysis of overrides.

The paper is organised as follows: In section 2 we discuss the connection between rating error rates and the statistical notion of misclassification rate. We argue that the misclassification rate related to a particular classification rule provides a suitable approximate upper bound for the rating error rate and therefore also for the rate of overrides of the outputs of a statistical rating model. In section 3 we describe how to calculate the override rate bounds in the special cases of a rating model with a finite number of grades and a rating model with normally distributed conditional score distributions. The case of the conditional normal distributions, in particular, allows us to study the connection between override rate bounds and discriminatory power. We also present some numerical examples to illustrate this connection. In section 4 we discuss how the proposed bound for rating override rates fits into an illustrative framework for the monitoring of rating overrides. Section 5 summarises and concludes the paper.

A note on the terminology. In this paper, we study rating overrides in the context of rating systems with a small finite number k of grades³. For the purpose of the paper, it is assumed that a *statistical model* is used to determine *proposed ratings* which may then be confirmed or overridden by experts to give the *final ratings*. The direct output of the statistical model is called *score* and may be discrete with a large (compared to the number of rating grades k) range of scores or on a continuous scale. The range of the scores is decomposed into a number of disjoint intervals each of which is then mapped to one of the rating grades, thus generating the proposed ratings. The decomposition of the score range for the mapping to the grades usually is based on the probabilities of default associated with the scores.

³In order to achieve comparability of their internal ratings to agency ratings, financial institutions sometimes chose k as seven, the number of performing unmodified grades used by the major rating agencies, or as seventeen, the number of performing modified grades used by the agencies. In the latter case, grade 17 would correspond to S&P grade AAA, grade 16 to S&P grade AA+ etc. until grade 1 that would correspond to S&P grade CCC.

Figure 1: Unconditional and conditional rating distributions. The distributions conditional on default and survival respectively have been inferred from the unconditional distribution under the assumption that the unconditional PD is 5% and the accuracy ratio associated with the rating model is either low at 25% or high at 75%. See example 3.3 for details of the calculations.



Convention. Low values of the model output score indicate low creditworthiness, high values of the score indicate high creditworthiness. The mapping of the scores to the rating grades (expressed as positive integers) is increasing, i.e. high grades mean high creditworthiness.

2 Approximating rating error rate by misclassification rate

In the following, we assume that for every real-life rating model there are unavoidable rating errors because defaults are unpredictable. We want to determine an approximation for the proportion of such unavoidable rating errors compared to all rating actions. We also assume that the rating models we consider are correctly calibrated such that no overrides are needed in order to account for PD estimation bias. For the purpose of this paper, such overrides would be considered avoidable (by recalibration of the rating model). Establishing bounds for unavoidable errors will help to identify overrides that are due to miscalibration and thus in principle avoidable.

2.1 The typical process for rating overrides

Let us look at a – typically wholesale (e.g. corporates or financial institutions) – portfolio of borrowers that are in the scope of application of the rating system under consideration. A rating action for one of the borrowers consists of three steps:

- For the first step, a statistical model is applied to data (*risk factors*) related to the borrower, typically financial ratios from the borrower’s balance sheet and/or qualitative assessments by credit officers marked up on a standard scale. The result of this step is a score s .
- For the second step, the score s is mapped (e.g. by means of a look-up table) to a grade $g^* \in \{1, \dots, k\}$. The grade g^* is called *proposed rating*.
- In the third step, the rating proposal g^* is reviewed by a credit expert or a committee of credit experts. The experts can decide to accept the rating proposal and assign the borrower the *final rating* $g = g^*$. But the experts can also decide to reject the rating proposal and assign the borrower a final rating $g \neq g^*$ (but with $g \in \{1, \dots, k\}$). If a proposed rating is rejected, the experts have to record both the proposed and the final ratings and the rationale for the rejection and the choice of the final rating.

The occurrence of $g \neq g^*$ in step three is called *override* (of the proposed rating). In most credit institutions overrides are subject to certain restrictions which could include the following:

- Overrides are not allowed if the proposed ratings are better than or equal to some threshold grade $k^* < k$. This rule might be established when it is felt that the relative differences between the better rating grades are so small that the grades cannot be meaningfully differentiated by the human mind.
- Overrides are only allowed if the final ratings are significantly different to the proposed ratings, e.g. if $|g - g^*| \geq b$ for some threshold value b . Such a rule might again be based on the intuition that the human mind cannot differentiate between adjacent rating grades.
- Only downgrade overrides are allowed. This rule would express a desire for rating conservatism.

If n is the number of rating actions⁴ with the rating system under consideration within a pre-defined time period (e.g. one year), resulting in pairs $(g_1^*, g_1), \dots, (g_n^*, g_n)$ of proposed and final ratings, then the *override rate* is defined as the ratio of the number of rating actions with different proposed and final ratings and the total number of rating actions:

$$\text{Override rate} = \frac{\#\{i : g_i^* \neq g_i\}}{n}. \quad (2.1)$$

2.2 ‘Right’ and ‘wrong’ ratings

From (2.1) we get the obvious upper bound of 100% for override rates – not very satisfactory. It might then seem a natural idea to try to derive an “expected” override rate for a statistical

⁴The same borrower might be rated several times but usually each borrower needs to be rated at least once per year.

rating model by analysing the decision processes of the credit experts charged with reviewing the rating proposals. One could argue that credit experts are likely to base their override decisions on comparisons to borrowers rated in the past. This would mean that the experts look for similarities between the borrower to be reviewed and the borrowers rated earlier. Such similarities might relate to financial conditions, management quality or other descriptive information about the borrowers.

Hence the procedure by credit experts to decide about possible overrides might be described in mathematical terms as a methodology similar to the “ m nearest neighbours” (m -NN) methodology (see, e.g., [Hand, 1997](#), section 5.3), which is well-known in fields like Pattern Recognition or Statistical Classification. However, this analogy is not a promising path to the prediction of override rates. Assuming we were able to identify the number m of neighbours to be included in the m -NN analysis, we would end up with developing another rating system especially for the override decision process. But there is no reason to be sure that such a competitor rating would be more predictive of creditworthiness than the statistical model under review.

A more promising approach to bounds for override rates is to start from an interpretation of overrides as corrective actions. Then all ‘wrong’ – in a sense to be determined – ratings proposed by the statistical model should be overridden. Hence the proportion of wrong proposed ratings would be the ‘natural’ override rate. As in practice it is impossible to identify wrong ratings with certainty the natural override rate would be an upper bound for the number of overrides. This follows from the following considerations:

- Assume the observed override rate exceeds the natural override rate. Then there are two possible explanations for this. The credit experts could have overridden too many ratings and, hence, have turned some correct ratings into wrong ratings. Or the proportion of wrong ratings is higher than expected and, hence, the discriminatory power of the statistical model is lower than predicted. Both explanations imply a deficiency in the rating process that should be rectified.
- Assume the observed override rate is less than the natural override rate. Again there are two possible explanations. The credit experts could have overridden too few ratings and, hence, have failed to amend all wrong ratings. Or the proportion of wrong ratings is lower than expected and, hence, the discriminatory power of the statistical model is higher than predicted. While the first explanation may or may not indicate a deficiency in the rating process, the second explanation actually would be good news.
- With a view on the potential consequences of these two cases (“override rate exceeds the natural override rate” and “override rate is below the natural override rate”), we observe that the first case is clearly more of a concern than the second case. For the first case indicates with certainty a deficiency of the rating process while in the second case there is a possibility that the statistical model actually is better than expected. This is why in practice the natural override rate can only be regarded as an upper bound for the observed override rate whose breach should trigger corrective action.

So far, so well – but how do we know which proportion of rating proposals is wrong? Clearly, if we knew which would be the right – or maybe only the most predictive – ratings we would be done because then we could insert the right ratings as the final ratings g_i in (2.1). Interestingly

enough, if we assume that the most predictive ratings can be determined based on the analysis of a finite number of risk factors, then in theory a most predictive statistical rating model exists. This follows from the Neyman-Pearson lemma (see, e.g., [Casella and Berger, 2002](#), Theorem 8.3.12) that identifies the most predictive statistic for the test of two simple hypotheses as the ratio of the two multi-variate joint densities of the risk factors on the defaulter population and the survivor population respectively. Unfortunately, even if all relevant risk factors could be identified, in practice it would be impossible to estimate accurately enough the two risk factor densities.

Alternative candidates for the most predictive ratings might be agency ratings as well as the results of expert ranking exercises. However, both agency ratings and expert rankings are not always available. More importantly, even if the quality of agency ratings for corporates is good in general, there is no reason to consider them the ‘right’ ratings, implying the ‘right’ ranking. If for nothing else, this follows from the fact that ratings from different agencies for the same entities coincide often but not always. By exception, the assessments by different agencies may differ significantly. Similarly, for samples of borrowers ranked by experts, the resulting rankings will depend more often than not on the selection of the experts and possibly also on the way the assessments by the single experts are combined to one ranking.

So, as it proves difficult to compare the statistical rating model in question to ‘right’ or most predictive ratings, why not comparing it to a perfect (or ‘prophetic’) rating system? This approach might seem strange at first glance, but it turns out to be viable – even if it delivers only an approximation to the natural override rate.

A perfect rating system obviously needs only two grades: Default and survival. We have to define a mapping from the rating grades that are assigned by the statistical rating model in question to the two states of default and survival in order to be able to calculate which portion of the borrowers in the portfolio would have to be moved were the distribution of borrowers across the range of proposed grades to be transformed into a distribution according to a perfect rating system. For this purpose, we assume that the perfect rating system for the comparison is ‘realised’ on the rating scale of the rating model under consideration. This implies that there is a multitude of perfect rating systems that could be used as the target for the comparison. For each rating system that assigns defaulters and survivors to disjoint sets of grades is perfect.

To devise the mapping from the proposed ratings to the perfect ratings, inspection of the concept of investment and speculative grades employed by the major rating agencies proves helpful. Moody’s define *Investment grade* as the combination of grades Aaa, Aa, A, and Baa while they call *Speculative grade* the combination of grades Ba, B, and Caa-C. Similarly, S&P define *Investment grade* as the combination of grades AAA, AA, A, and BBB and *Speculative grade* as the combination of grades BB, B, and CCC-C.

Table 1 presents long-run average annual default rates for investment and speculative grade borrowers as observed by Moody’s and S&P. On the one hand, these default rates show that, from a credit default perspective, Moody’s and S&P’s investment grade and speculative grade definitions are broadly equivalent. On the other hand, and more important for this paper, the observed default rates support the description of investment grade borrowers as ‘very unlikely to default within a year’ and of speculative grade borrowers as ‘at significant risk of default within

Table 1: Average annual default rates as recorded by **Moody's** (2011, Exhibit 35) and **S&P** (2011, Table 24).

Agency	Investment grade	Speculative grade	Observation period
Moody's	0.095%	4.944%	1983-2010
S&P	0.13%	4.36%	1981-2010

a year'. One might even be tempted to talk about investment grade borrowers as 'safe' borrowers and about speculative grade borrowers as 'risky' borrowers. This can be seen as a reasonable approximation of a perfect rating system with the two grades 'defaulter' and 'survivor' only.

Intuitively, a coarse classification of borrowers into safe and risky should be easier to achieve than the much finer differentiation envisaged by the rating agencies and most internal rating systems of financial institutions. Indeed, the restrictions of overrides discussed in section 2.1 might be rationalised by the consideration that it is primarily the safe vs. risky differentiation that one has to get right.

We adopt this view for our approach to the identification of bounds for override rates with 'natural' error rates. Hence, we will define the natural error rate as the misclassification rate in a two-state coarse rating system that is derived from the statistical rating model in question by combining a suitable subset of grades to a 'safe' super-grade and the complementary subset of grades to a 'risky' super-grade.

2.3 Natural error rate

We consider first how to determine the sets of scores or grades that define the two super-grades safe and risky. Then we discuss how to calculate the corresponding misclassification rate that will be defined as the natural error rate associated with the statistical rating model in question.

Actually, these two problems can be solved in a fully general context. Hence although in principle we only need to investigate the cases of a rating variable with a finite range of values and of a continuous real-valued score variable, the major part of the following discussion applies to any random variable S with values in a general measurable space – only minor modifications of the notation might be necessary. In practice the measurable space will be a multi-dimensional Euclidian space \mathbb{R}^n when we are dealing with the risk factors informing a rating model, an open real interval $I \subset \mathbb{R}$ when we are talking about the score output of a rating model, or a finite ordered set (without loss of generality $\{1, 2, \dots, k\}$, with k being the number of performing rating grades) when the discussion is about the ratings proposed by a rating model or the final ratings after overrides.

Notation. We assume that the conditional and unconditional distributions of S have densities with respect to a suitable measure. To make it clear when a statement or an equation is general we adopt the *likelihood function* notation where a likelihood ℓ can stand both for a *probability density function* in a continuous or general context or for a *probability mass function* in a discrete

context.

Speaking in technical terms, in this paper we study the joint distribution of a pair (S, Z) of random variables. As mentioned before, the meaning of the variable S could be a vector of risk factors, a *rating model score*, or a *rating grade*, observed for a solvent borrower at a fixed date. The variable Z is the *borrower's state of solvency* one observation period (usually one year) after S was observed. Z takes on the two values 0 and 1. The meaning of $Z = 0$ is ‘borrower has remained solvent’ (solvency or survival), $Z = 1$ means ‘borrower has become insolvent’ (default). We write D for the event $\{Z = 1\}$ and N for the event $\{Z = 0\}$. Hence

$$D \cap N = \{Z = 1\} \cap \{Z = 0\} = \emptyset, \quad D \cup N = \text{whole space}. \quad (2.2)$$

The marginal distribution of the state variable Z is characterised by the *unconditional probability of default* p which is defined as the probability of the default event D :

$$p = P[D] = P[Z = 1] \in (0, 1). \quad (2.3)$$

The joint distribution of (S, Z) then can be specified by p and the two distributions of S conditional on the states of Z (i.e. the events D and N respectively). Most important are the cases where the conditional distributions are given by discrete conditional probabilities $\ell_D(s) = P[S = s | D]$ and $\ell_N(s) = P[S = s | N]$, $s = 1, \dots, k$, or by Lebesgue densities ℓ_D and ℓ_N . In the latter case, the probabilities of S taking a value in a set M conditional on default and survival respectively can be written as

$$P[S \in M | D] = \int_M \ell_D(s) ds \quad \text{and} \quad P[S \in M | N] = \int_M \ell_N(s) ds. \quad (2.4)$$

The joint distribution of the pair (S, Z) of the score and the borrower's state one period later can also be specified by the unconditional distribution $P[S \in \cdot]$ of S (i.e. the score or rating profile) and the PDs $P[D | S] = 1 - P[N | S]$ *conditional* on S . In the special cases we have mentioned before, by Bayes' rule we have the following representations of the conditional PDs:

- S is understood as rating grade, i.e. $S \in \{1, 2, \dots, k\}$. Then

$$P[D | S = s] = \frac{p P[S = s | D]}{p P[S = s | D] + (1 - p) P[S = s | N]}, \quad s \in \{1, 2, \dots, k\}. \quad (2.5a)$$

- S is a continuous score variable or a vector of continuous risk factors such that there are Lebesgue densities ℓ_N and ℓ_D of the distributions conditional on default and survival as in (2.4). Then

$$P[D | S = s] = \frac{p f_D(s)}{p f_D(s) + (1 - p) f_N(s)}. \quad (2.5b)$$

With the likelihood notation, (2.5a) and (2.5b) can be expressed in one equation:

$$P[D | S = s] = \frac{p \ell_D(s)}{p \ell_D(s) + (1 - p) \ell_N(s)}. \quad (2.5c)$$

Thanks to (2.5c) the conditional PDs are determined by the unconditional PD and the distributions of the score conditional on default and survival. The conditional and the unconditional

PDs are the information we need to determine the intended split of the score or rating range into the two super-grades safe and risky. The split suggested is based on the minimisation of the so-called *expected misclassification cost* (see, e.g., [Hand and Till, 2001](#)).

Proposition 2.1 *Assume that the expected cost of misclassifying a defaulting borrower as solvent is $c_D > 0$ and that the expected cost of misclassifying a solvent borrower as defaulting is $c_N > 0$. Define the events D and N as in (2.2) and the unconditional PD p as in (2.3). Assume that a borrower is classified as defaulting if and only if the borrower's score (or rating or vector of risk factors) S takes on a value in a predefined set A . Then the expected cost C of misclassification is given by*

$$C = c_D p P[S \notin A | D] + c_N (1 - p) P[S \in A | N]. \quad (2.6)$$

The expected cost C of misclassification is minimal if and only if the set A is chosen as the following set A^* :

$$A^* = \left\{ P[D | S] > \frac{c_N}{c_N + c_D} \right\}. \quad (2.7)$$

See, for instance, [Nechval et al. \(2010, Theorem 2\)](#) for a proof of proposition 2.1. Note that (2.7) defines a risky super-grade in the sense that the average PD conditional on the score or rating on the set A^* is greater than the average PD conditional on the score or rating on the complement of the set A^* :

$$E[P[D | S] | S \in A^*] > \frac{c_N}{c_N + c_D} \geq E[P[D | S] | S \notin A^*]. \quad (2.8)$$

How to choose the cost parameters c_D and c_N in (2.7)? [Hand \(2009, section 4\)](#) has commented that in general it is extremely difficult to determine values for the cost parameters. Even as it obviously suffices to specify the ratio of the cost parameters, most of time it appears difficult to come to a meaningful assessment.

Nonetheless, as noted by [Hand \(2009, section 4\)](#), choosing the cost parameters inversely proportional to the unconditional class probabilities is popular. In the setting of this paper, this means

$$c_D = 1/p \quad \text{and} \quad c_N = 1/(1 - p). \quad (2.9)$$

In a context of risky lending, (2.9) is actually a quite reasonable choice. Assume that the credit decision is about a loan with principal M and that the expected loss rate in the case of default is $0 < \lambda \leq 1$. Then the effective (after cost of funding) interest charged for the loan should be at least a provision for the expected loss which equals $M \lambda p$. If we prudentially assume that the interest is paid in advance then the lending institution's loss in the case of lending to a defaulter will be

$$M \lambda - M \lambda p = M \lambda (1 - p). \quad (2.10a)$$

The loss in case of not lending to a customer who turns out not to be a defaulter will be (at least) $M \lambda p$. Together with (2.10a) this implies for the ratio of the misclassification cost parameters

$$\frac{M \lambda (1 - p)}{M \lambda p} = \frac{1/p}{1/(1 - p)} = \frac{c_D}{c_N}, \quad (2.10b)$$

with c_D and c_N as suggested in (2.9). Even if the assumption on interest being paid in advance clearly is often not satisfied in practice, for not too large PD p the ratio of misclassification costs will still be of a magnitude similar to the left-hand side of (2.10b). We therefore adopt (2.9) for the choice of the cost parameters in our application of proposition 2.1.

Observe that with cost parameters as defined by (2.9), by equations (2.5c) and (2.7) the optimal set A^* of grades or ratings for classifying borrowers as defaulting can be rewritten as follows:

$$\begin{aligned} A^* &= \left\{ \frac{1}{1 + ((1-p)/p) (\ell_N(S)/\ell_D(S))} > \frac{1}{1 + c_D/c_N} \right\} \\ &= \left\{ \frac{\ell_D(S)}{\ell_N(S)} > \frac{1-p}{p} \frac{c_N}{c_D} \right\} \end{aligned} \quad (2.11a)$$

$$= \left\{ \frac{\ell_D(S)}{\ell_N(S)} > 1 \right\}. \quad (2.11b)$$

Comparison of (2.11a) and (2.11b) shows that the same optimal set A^* for classifying borrowers as defaulting is found, no matter whether we start from a general unconditional PD p and cost parameters given by (2.9) or from $p = 50\%$ and equal cost parameters $c_D = c_N$.

Remark 2.2 Assume that the variable S from proposition 2.1 stands for a one-dimensional score or rating. If the likelihood ratio $\frac{\ell_D(s)}{\ell_N(s)}$ is decreasing⁵ in s (or, equivalently, the PD conditional on the score $P[D|S = s]$ is increasing in s) then the set A^* from (2.7) indicating the classification of borrowers as defaulting can be written as

$$\begin{aligned} A^* &= \{S \leq s^*\}, \\ s^* &= \sup \left\{ s : \frac{\ell_D(s)}{\ell_N(s)} > 1 \right\}. \end{aligned} \quad (2.12)$$

In this case the minimum classification cost C^* according to (2.6) with cost parameters defined by (2.9) is equivalent to the Kolmogorov-Smirnov statistic applied to the distribution functions of the score or rating conditional on default and survival respectively:

$$C^* = 1 - \max_s |P[S \leq s | D] - P[S \leq s | N]|. \quad (2.13)$$

This is an interesting observation since it justifies, in economic terms, the use of the Kolmogorov-Smirnov statistic and of the closely related accuracy ratio for the performance measurement of rating systems.

With (2.11b) we have identified the decision rule that should be applied to determine the two-state coarse rating system that was discussed in section 2.2. This rating system is an overlay to the statistical rating model whose output is the variable S we have considered in proposition 2.1 and its conclusions. The misclassification rate of the overlaying rating system with the two super-grades ‘risky’ and ‘safe’ gives us a reasonable bound for the override rate of the statistical rating model S .

⁵See Tasche (2008) for a discussion of the connection between monotonicity of the likelihood ratio and optimality of the model scores.

Proposition 2.3 *Define the natural error rate $\epsilon(S)$ of a statistical rating model with output S as the misclassification rate associated with the decision rule (2.11b) for classifying borrowers as defaulting. Let p denote the unconditional probability of default. Then $\epsilon(S)$ is given by*

$$\epsilon(S) = p \mathbb{P}[\ell_D(S) \leq \ell_N(S) | D] + (1 - p) \mathbb{P}[\ell_D(S) > \ell_N(S) | N]. \quad (2.14a)$$

Note that the natural error rate as defined by (2.14a) is an *actual error rate* in the terminology of Hand (1997, section 7.2). If the output of the statistical rating model is a rating proposal or a discrete score with values in $\{1, \dots, k\}$ then the likelihood functions in (2.14a) are probability mass functions. Let

$$J = \{j : \mathbb{P}[S = j | D] > \mathbb{P}[S = j | N]\}. \quad (2.14b)$$

Then (2.14a) can be written as

$$\epsilon(S) = p \sum_{j \notin J} \mathbb{P}[S = j | D] + (1 - p) \sum_{j \in J} \mathbb{P}[S = j | N]. \quad (2.14c)$$

If the output of the statistical rating model is a continuous score with values in an open interval I then the likelihood functions in (2.14a) are probability density functions. (2.14a) can then be represented as

$$\epsilon(S) = p \int_{\{s \in I : \ell_D(s) \leq \ell_N(s)\}} \ell_D(s) ds + (1 - p) \int_{\{s \in I : \ell_D(s) > \ell_N(s)\}} \ell_N(s) ds. \quad (2.14d)$$

As mentioned before, the purpose of this paper is to demonstrate that the natural error rate of a statistical rating model as defined formally in proposition 2.3 is a suitable bound for the rate of overrides that should be applied to the model's rating proposals. In section 3, we show that the natural error rate has the intuitive property that its value increases when the discriminatory power of the statistical rating models declines. This observation strengthens the case we have made in section 2.2 for the adoption of the natural error rate as a bound for the rating override rate.

It is worthwhile noting, however, that a regrouping of borrowers according to the decision rule (2.11b) underlying the coarse approximate rating system would imply moving a much larger portion of solvent borrowers from the 'risky' grade to the 'safe' grade than of defaulting borrowers from safe to risky. This observation might seem counterintuitive at first glance. Nonetheless, in practice the movement from risky to safe would be quite restricted as only borrowers that have been very clearly misclassified should be moved. This consideration indicates that the 'natural error rate' indeed is an upper bound for the rating override rate since most of the misclassified borrowers cannot be identified with reasonable certainty.

3 Override rate and discriminatory power

Define the conditional distribution functions F_D and F_N of the score or rating variable S described in section 2.3 by

$$F_D(s) = \mathbb{P}[S \leq s | D] \quad \text{and} \quad F_N(s) = \mathbb{P}[S \leq s | N], \quad (3.1)$$

and denote by S_D and S_N independent random variables that are distributed according to F_D and F_N respectively.

For quantifying discriminatory power, we apply the notion of *accuracy ratio* (AR) as specified in Tasche (2009, eq. (3.28b)):

$$\begin{aligned} AR &= 2\mathbb{P}[S_D < S_N] + \mathbb{P}[S_D = S_N] - 1 \\ &= \mathbb{P}[S_D < S_N] - \mathbb{P}[S_D > S_N] \\ &= \int \mathbb{P}[S < s | D] dF_N(s) - \int \mathbb{P}[S < s | N] dF_D(s), \end{aligned} \tag{3.2}$$

See Hand and Till (2001, section 2) for a discussion of why this definition of accuracy ratio (or the related definition of the area under the ROC curve) is more expedient than the also common definition in geometric terms. Definition (3.2) of AR takes an ‘ex post’ perspective by assuming the obligors’ states D or N one year after having been scored are known and hence can be used for estimating the conditional (on default and survival respectively) score distributions F_D and F_N .

3.1 The binormal case

The ex-post perspective on discriminatory power is appropriate when we consider the important binormal special case where both conditional score distributions are normal distributions with different mean values but equal variances. This special case, in particular, motivates the choice of the inverse logit function for modelling PD curves (see, e.g., Cramer, 2003, section 6.1) since in practice it is often reasonable to assume that the two conditional score distributions are approximately normal.

Proposition 3.1 *Assume that the conditional distributions of the score variable S are given by $S_D \sim \mathcal{N}(\mu_D, \sigma)$ and $S_N \sim \mathcal{N}(\mu_N, \sigma)$ with $\mu_D \leq \mu_N$. Let p denote the unconditional probability of default as given by (2.3). Then the accuracy ratio AR associated with S according to (3.2) can be calculated as⁶*

$$AR = 2\Phi\left(\frac{\mu_N - \mu_D}{\sqrt{2}\sigma}\right) - 1. \tag{3.3a}$$

The natural error rate $\epsilon(S)$ of S according to proposition 2.3 is given by

$$\epsilon(S) = \Phi\left(\frac{\mu_D - \mu_N}{2\sigma}\right). \tag{3.3b}$$

With the cost parameters defined by (2.9), the average conditional PDs on the set of scores indicating default and on the set of scores indicating survival respectively (as in eq. (2.8)) are given by

$$\begin{aligned} \mathbb{E}[\mathbb{P}[D | S] | S \in A^*] &= \frac{p(1 - \epsilon(S))}{p(1 - \epsilon(S)) + (1 - p)\epsilon(S)}, \\ \mathbb{E}[\mathbb{P}[D | S] | S \notin A^*] &= \frac{p\epsilon(S)}{p\epsilon(S) + (1 - p)(1 - \epsilon(S))}. \end{aligned} \tag{3.3c}$$

⁶ Φ denotes the standard normal distribution function.

Proof. (3.3a) follows from Tasche (2009, eq. (3.14) and proposition 3.15). For (3.3b) and (3.3c), observe that in the binormal case with equal variances we have

$$S \notin A^* \iff \ell_D(S) \leq \ell_N(S) \iff S \geq \frac{\mu_D + \mu_N}{2}.$$

From this observation, we deduce that

$$\mathbb{P}[\ell_D(S) \leq \ell_N(S) \mid D] = 1 - \Phi\left(\frac{\mu_N - \mu_D}{2\sigma}\right)$$

and

$$\mathbb{P}[\ell_D(S) > \ell_N(S) \mid N] = \Phi\left(\frac{\mu_D - \mu_N}{2\sigma}\right).$$

These equations imply (3.3b). With regard to (3.3c), we apply the definition of conditional expectation to arrive at the following computation⁷:

$$\begin{aligned} \mathbb{E}[\mathbb{P}[D \mid S] \mid S \in A^*] &= \frac{\mathbb{P}[\mathbb{P}[D \mid S] \mathbf{1}_{\{S \in A^*\}}]}{\mathbb{P}[S \in A^*]} \\ &= \frac{\mathbb{P}[D \cap \{S < \frac{\mu_D + \mu_N}{2}\}]}{\mathbb{P}[S < \frac{\mu_D + \mu_N}{2}]} \\ &= \frac{\mathbb{P}[S < \frac{\mu_D + \mu_N}{2} \mid D] \mathbb{P}[D]}{\mathbb{P}[S < \frac{\mu_D + \mu_N}{2} \mid D] \mathbb{P}[D] + \mathbb{P}[S > \frac{\mu_D + \mu_N}{2} \mid N] \mathbb{P}[N]} \\ &= \frac{p \Phi\left(\frac{\mu_N - \mu_D}{2\sigma}\right)}{p \Phi\left(\frac{\mu_N - \mu_D}{2\sigma}\right) + (1 - p) \Phi\left(\frac{\mu_D - \mu_N}{2\sigma}\right)}. \end{aligned}$$


Since $\Phi\left(\frac{\mu_N - \mu_D}{2\sigma}\right) = 1 - \epsilon(S)$ we obtain the first equation in (3.3c). The second equation follows in the same way. \square

In proposition 3.1, it is a general property of the accuracy ratio that it does not depend on the unconditional PD. The observation that the natural error rate does not depend on the unconditional PD might be surprising at first glance. This observation, however, is primarily a consequence of the fact that both conditional score distributions are members of the same location-scale distribution family, different only by the different location parameters. In addition, symmetry of the location zero member of the family is required to imply the property that the error does not depend on the unconditional PD.

Combining (3.3a) and (3.3b) provides a simple relationship between the discriminatory power as measured by the accuracy ratio and the natural error rate in the binormal case with equal variances. Eq. (3.4) shows, in particular, that the natural rating error indeed decreases when the discriminatory power of the statistical rating model increases.

⁷We define the indicator function $\mathbf{1}_M$ of a set M by $\mathbf{1}_M(m) = \begin{cases} 1, & m \in M, \\ 0, & m \notin M. \end{cases}$

Figure 2: Natural error rate as function of the discriminatory power (accuracy ratio) for the binormal case described in corollary 3.2 and for the discrete rating model described in example 3.3.



Corollary 3.2 Assume that the conditional distributions of the score variable S are given by $S_D \sim \mathcal{N}(\mu_D, \sigma)$ and $S_N \sim \mathcal{N}(\mu_N, \sigma)$ with $\mu_D \leq \mu_N$. Then the accuracy ratio AR and the natural error rate $\epsilon(S)$ associated with S are related by the following equation:

$$\epsilon(S) = \Phi \left(-\frac{1}{\sqrt{2}} \Phi^{-1} \left(\frac{AR + 1}{2} \right) \right). \quad (3.4)$$

See figure 2 for a graphical representation of relation (3.4).

3.2 The discrete case

For practical applications, often one cannot assume that the conditional score distributions are normally distributed. Moreover, as described in section 2.1, it is actually not the score output of a statistical rating model that may be overridden but the proposed ratings derived from the scores. Typically, thresholds for override rates for a statistical rating model are established when the development of the model is being finished and the documentation of the related rating system is compiled. At this stage, it is unlikely that reliable estimates of the score distribution

conditional on the borrower's solvency state or of the proposed rating distributions conditional on the borrower's solvency state are available. Therefore, we also consider the case where the unconditional distribution of the proposed ratings (rating profile) and the associated PD curve (i.e. the estimates of the PDs conditional on the proposed ratings) are all that is known.

Nonetheless, with this information it is still possible to compute both the natural error rate and the accuracy ratio for the statistical rating model in question. Key to these computations is the following observation that eq. (2.5a) can be 'inverted' such that the distributions of the proposed ratings conditional on the borrower's solvency state can be derived from the unconditional rating profile and the PD curve:

$$\begin{aligned} p &= E[P[D | S]] = \sum_{s=1}^k P[D | S = s] P[S = s], \\ P[S = s | D] &= P[D | S = s] P[S = s] / p, \quad s = 1, \dots, k, \\ P[S = s | N] &= (1 - P[D | S = s]) P[S = s] / (1 - p), \quad s = 1, \dots, k. \end{aligned} \tag{3.5}$$

By (3.5) and (2.14c), the natural error rate can be calculated from the profile $s \mapsto P[S = s]$ of the proposed ratings and the PD curve $s \mapsto P[D | S = s]$. By combining (3.2) and (3.5), also the accuracy ratio of the ratings proposed by the statistical model can be determined from the rating profile and the PD curve:

$$\begin{aligned} AR &= \frac{1}{p(1-p)} \left(2 \sum_{s=1}^k (1 - P[D | S = s]) P[S = s] \sum_{t=1}^{s-1} P[D | S = t] P[S = t] \right. \\ &\quad \left. + \sum_{s=1}^k P[D | S = s] (1 - P[D | S = s]) P[S = s]^2 \right) - 1. \end{aligned} \tag{3.6}$$

Calculation of the accuracy ratio by means of the combination of (3.5) and (3.6) may be interpreted as predicting the accuracy ratio 'ex ante' since it can be done on the current portfolio as soon as estimates of the PDs conditional on the rating grades (PD curve) are available.


There is no obvious general example of a discrete rating model comparable to the binormal model from corollary 3.2 that would allow the direct study of the connection between discriminatory power and natural error rate. We therefore explore this connection in the discrete case by means of a specific example. In order to make the example differ significantly from the normal assumption discussed before it has been chosen in a way as to create a certain degree of unsymmetry and overdispersion.

Example 3.3 *The unconditional distribution of the ratings on a discrete scale from 1 to 17 is given by a correlated binomial distribution. That is if S stands for a borrower's rating grade, S can be written as $S = X + 1$ where the distribution of X is specified by⁸*

$$\begin{aligned} P[X \leq x] &= \int_{-\infty}^{\infty} \varphi(y) \sum_{i=0}^x \binom{k}{i} G(\lambda, \varrho, y)^i (1 - G(\lambda, \varrho, y))^{k-i} dy, \quad x = 0, \dots, k, \\ G(\lambda, \varrho, y) &= \Phi\left(\frac{\Phi^{-1}(\lambda) - \sqrt{\varrho} y}{\sqrt{1 - \varrho}}\right). \end{aligned} \tag{3.7a}$$

⁸ φ denotes the standard normal density function: $\varphi(t) = \frac{e^{-t^2/2}}{\sqrt{2\pi}}$.

Figure 3: PD curves for the rating model described in example 3.3. The curves have been inferred from the unconditional rating distribution shown in figure 1 under the assumption that the unconditional PD is 5% and the accuracy ratio associated with the rating model is either low at 25% or high at 75%.



For the purpose of this example we chose the following values for the parameters k (number of grades minus one), λ (determines mean of the distribution) and ϱ (drives overdispersion compared to binomial distribution) in (3.7a):

$$k = 16, \quad \lambda = 0.55, \quad \varrho = 0.1. \quad (3.7b)$$

We assume that the PD conditional on a rating grade $S = s$ is appropriately described by the inverse logit function:

$$P[D | S = s] = \frac{1}{1 + e^{a+bs}}, \quad s = 1, \dots, 17. \quad (3.8)$$

The parameters a and b are determined by quasi-moment matching (*Tasche, 2009*, section 5.2). This approach works by equating the right-hand side of the first equation of (3.5) and the right-hand side of (3.6) to predefined values of PD and AR respectively and solving numerically for

Table 2: Natural error rate as function of the discriminatory power (accuracy ratio) for the binormal case described in corollary 3.2 and for the discrete rating model described in example 3.3. See figure 2 for a graphical representation.

AR	Natural error rate		
	Binormal	Discrete, PD=1%	Discrete, PD=10%
0.0	0.500	0.452	0.461
0.1	0.465	0.451	0.448
0.2	0.429	0.450	0.435
0.3	0.393	0.448	0.422
0.4	0.355	0.320	0.306
0.5	0.317	0.319	0.291
0.6	0.276	0.317	0.276
0.7	0.232	0.206	0.260
0.8	0.182	0.204	0.154
0.9	0.122	0.118	0.131

a and b :

$$\begin{aligned}
 PD &= \sum_{s=1}^{17} P[S = s] P[D | S = s] = \sum_{s=1}^{17} \frac{P[X = s - 1]}{1 + e^{a+bs}}, \\
 AR &= \frac{1}{PD(1 - PD)} \left(2 \sum_{s=1}^{17} \frac{e^{a+bs}}{1 + e^{a+bs}} P[X = s - 1] \sum_{t=1}^{s-1} \frac{P[X = t - 1]}{1 + e^{a+bt}} \right. \\
 &\quad \left. + \sum_{s=1}^{17} \frac{e^{a+bs}}{(1 + e^{a+bs})^2} P[X = s - 1]^2 \right) - 1.
 \end{aligned} \tag{3.9}$$

Figure 3 illustrates the results of quasi-moment matching according to (3.9). In particular, it becomes clear that the slope of the PD curve is primarily controlled by the discriminatory power of the rating model as expressed by its accuracy ratio. The two curves from figure 3 have been used together with the unconditional rating distribution as specified by (3.7a) to calculate (by means of (3.5)) the conditional rating distributions that are shown in figure 1. Once the conditional rating distributions have been determined, the natural error rates (and hence bounds for the overrides) may be computed by means of equations (2.14b) and (2.14c).

3.3 Comparing the binormal and the discrete results

The same approach that was used for the calculation of the conditional rating distributions in figure 1 has also been applied in order to calculate the discrete error rate curves for figure 2, albeit with different parameters. For figure 2, the calculations were done for a small unconditional PD (1%) and a large unconditional PD (10%). Each of the two PDs was combined with the whole range $[0, 1]$ of potential ARs to first determine conditional PD curves by means of (3.9) and

Table 3: Average ‘safe’ and ‘risky’ PDs as function of the discriminatory power (AR) for the binormal case (3.3c) and example 3.3. The assumed unconditional PD is 1%.

AR (%)	Binormal		Discrete	
	Safe PD (%)	Risky PD (%)	Safe PD (%)	Risky PD (%)
0	1	1	1	1
10	0.87	1.15	0.87	1.16
20	0.75	1.33	0.74	1.32
30	0.65	1.54	0.61	1.47
40	0.55	1.8	0.57	1.9
50	0.47	2.13	0.45	2.15
60	0.38	2.58	0.34	2.39
70	0.3	3.24	0.32	3.54
80	0.22	4.33	0.19	4.02
90	0.14	6.75	0.13	7.06

then natural error rates by means of (2.14b) and (2.14c). The discontinuities in the ‘discrete’ curves are owed to changes in the index set J from (2.14b) which cause jumps in the values of the error rates. Despite the jumps, the two discrete curves are remarkably close to the binormal curve from corollary 3.2. This observation is confirmed by table 2 that provides the numerical values for some of the points on the three curves in figure 2. It appears therefore worthwhile to consider that the ‘normal’ values calculated with (3.4) are taken to indicate the link between discriminatory power and override even for non-normal cases.

In section 2.2, the approach to override rate bounds via the misclassification rate of a coarse two-state rating system was motivated by an inspection of the investment and speculative grades concept used by the major rating agencies. By table 1, we had noted that investment grade may be interpreted as ‘safe’ and speculative grade may be regarded as ‘risky’. Table 3 presents ‘risky’ and ‘safe’ grade PDs in the sense of equation (2.8) for the normal and discrete examples considered in this section.

Note the similarity between the average PDs from the 80% accuracy ratio row of table 3 and the observed investment and speculative grade default rates from table 1. Moody’s report an overall average annual default rate of 1.8% (Moody’s, 2011, Exhibit 35) and an average one-year accuracy ratio of c. 85% (Moody’s, 2011, Exhibit 15). S&P report an overall average annual default rate of 1.6% (S&P, 2011, Table 24) and an average one-year accuracy ratio of c. 84% (S&P, 2011, Table 2). Hence the observed similarity between average PDs on the ‘risky’ and ‘safe’ super-grades and the investment and speculative grade default rates recorded by the agencies might not be an incident. Rather one might guess that there is an expected cost concept like the one presented in proposition 2.1 behind the agencies’ investment versus speculative classification.

4 Monitoring rating overrides

In order to put the considerations from the previous sections into context, in this section we describe an illustrative framework for the monitoring of rating overrides. In particular, we make suggestions for the role that the natural error rate defined in proposition 2.3 could play as a bound (limit) for the override rate in such a framework.

Let us start by recalling the observations on statistical rating models and rating overrides that have lead us to suggesting the natural error rate as possible limit for rating override rates.

- No statistical rating model can predict credit defaults with certainty. In particular if the financial consequences of wrong rating decisions can be serious, therefore, rating proposals by models should be scrutinised by credit experts and be overridden if necessary.
- However, also the credit experts can be wrong with their decisions. Hence, there is a need to monitor the performance of the overrides, too.
- On a single case basis, it is not possible to decide whether a rating decision was right or wrong. Even an override of a CCC rating proposal to a AAA rating, followed immediately by the borrower's default, could in principle be just an occurrence of bad luck.
- Nonetheless, on the basis of the expected cost of misclassification, it is possible to identify sets of 'safe' and 'risky' rating grades, thus overlaying a 'coarse' rating system onto the original system with a finer rating scale. The safe and risky super-grades represent a 'will survive' and a 'will default' prediction respectively.
- The expected error rate of the 'coarse' rating system is called *natural error rate*. If the application of overrides is restricted to fundamental errors (i.e. safe borrowers are classified as risky or risk borrowers are classified as safe) the natural error rate is also a *natural override rate*.
- On portfolios with sufficient numbers of default observations, the natural error rate can be directly estimated (ex post). On portfolios with low numbers of default observations, the natural error rate can still be inferred from the rating profile and the PD curve of the rating model (ex ante).
- An observed override rate higher than the natural error rate could be a consequence of unnecessary and potentially erroneous overrides or of an overestimation of the discriminatory power of the rating model. In either case, the rating process would have been found to operate sub-optimally.
- An observed override rate lower than the natural error rate could be a consequence of overlooked rating errors or of an underestimation of the discriminatory power of the rating model. No clear immediate conclusion on the well- or malfunction of the rating process would be possible.

Based on these observations and with a view on affirming correct calibration of the rating model, an overriding monitoring framework could include the following elements:

At the beginning of the observation period (typically one year)

- Infer the natural error rate from the rating profile and the PD curve, by making use of

(2.14c) and (3.5), or from the accuracy ratio, by (3.4).

- Discourage credit experts from minor rating override. Only fundamental rating errors should be amended by overrides.

At the end of the observation period

- If there is a sufficient number of default observations, the discriminatory power of the rating system pre and post overrides should be computed. If the discriminatory power post overrides is significantly lower than the power pre overrides an update of the override governance might be necessary.
- If the observed override rate exceeds the natural error rate, the consistency of the PD curve with the observed discriminatory power should be checked (if there is a sufficient number of defaults) by comparing it to the ex ante accuracy ratio (3.6) because the slope of the PD curve might be too large. If the number of observed defaults is not sufficient for the ex post estimation of discriminatory power it should be confirmed that the overrides were justified (e.g. by interviews with the credit experts). If this is confirmed the PD curve must be amended as to reflect the lower than expected discriminatory power.
- If the observed override rate is lower than the natural error rate it should be confirmed (e.g. by internal audit) that the due procedure is followed with regard to overrides. If this is the case amendment of the PD curve as to reflect the higher than expected discriminatory power could be considered.
- An imbalance of upward and downward rating overrides with more upward overrides might be an indication of PD overestimation. If further analysis finds (e.g. based on statistical tests pre and post overrides if there is a sufficient number of default observations) that there is no overestimation an update of the override governance might be necessary.
- An imbalance of upward and downward rating overrides with more downward overrides might be an indication of PD underestimation. If further analysis finds that there is no underestimation an update of the override governance might be necessary.

5 Conclusions

The first part of this paper presents a suggestion of how to determine bounds for the rates of overrides of rating proposals determined by statistical rating models. In practice, often the guidance for such overrides tends to restrict them to essential cases where the model's assessments of borrowers as risky or safe might not be correct. Motivated by this observation, the suggested bound for the override rate is the misclassification rate of a coarse rating system with the two grades 'risky' and 'safe'. This coarse rating system is fed by the output of the rating model in question and combines its rating grades to two super-grades, thereby minimising the expected cost of misclassification. The rate of misclassifications with the coarse rating system is called 'natural error rate'. It is argued that the natural error rate is an appropriate bound for the override rate of the statistical rating model in question.

In the second part of the paper, methods for determining the natural error rate and its use in

an illustrative framework for the monitoring of overrides are discussed. It turns out that there is a particularly simple formula (eq. (3.4)) for the natural error rate if the rating distributions conditional on the borrower’s solvency state are both normal with equal variance (binormal case). Formula (3.4) for the natural error rate suggests that the natural error rate decreases when the discriminatory power of the rating model increases – as it should intuitively be expected. We compare the results from the ‘binormal’ formula to results for the natural error rate from a more realistic discrete-valued example of a rating model. The comparison shows the results by both approaches to be very close. This observation indicates that the binormal formula might be used in general as a rule of thumb for deriving a bound for the override rate from the discriminatory power of the rating model.

As demonstrated in the paper, the natural error rate of a rating model can be calculated both ex post, based on the discriminatory power realised during a previous observation period, and ex ante at the beginning of the observation period. In the latter case one applies Bayes’ rule to determine the discriminatory power implied by the combination of rating distribution and PD curve associated with the rating model. Both the ex post and the ex ante methods to the computation of the natural error rate as a bound for the override rate are promising tools for validating and monitoring the performance of statistical rating models.

References

- BCBS. *International Convergence of Capital Measurement and Capital Standards. A Revised Framework, Comprehensive Version*. Basel Committee on Banking Supervision, June 2006.
- M. Brown, M. Schaller, S. Westerfeld, and M. Heusler. Information or insurance? On the role of discretion in relationship lending. Working paper, University of St. Gallen, 2012.
- G. Casella and R.L. Berger. *Statistical Inference*. Duxbury Press, second edition, 2002.
- J.S. Cramer. *Logit Models From Economics and Other Fields*. Cambridge University Press, 2003.
- D.J. Hand. *Construction and Assessment of Classification Rules*. John Wiley & Sons, Chichester, 1997.
- D.J. Hand. Measuring classifier performance: a coherent alternative to the area under the ROC curve. *Machine Learning*, 77(1):103–123, 2009.
- D.J. Hand and R.J. Till. A simple generalisation of the area under the ROC curve for multiple class classification problems. *Machine Learning*, 45(2):171–186, 2001.
- Moody’s. Corporate Default and Recovery Rates, 1920-2010. Special comment, Moody’s Investors Service, February 2011.
- N.A. Nechval, K.N. Nechval, M. Purgailis, U. Rozevskis, V.F. Strelchonok, J. Krasts, G. Berzins, and K. Cikste. Modified approaches to pattern recognition via the Fisher discriminant ratio. In *Proceedings of the 10th International Conference “Reliability and Statistics in Transportation and Communication”*, pages 192–198, 2010.

- S&P. Default, Transition, and Recovery: 2010 Annual Global Corporate Default Study And Rating Transitions. Report, Standard & Poor's, March 2011.
- D. Tasche. Validation of internal rating systems and PD estimates. In G. Christodoulakis and S. Satchell, editors, *The Analytics of Risk Model Validation*, pages 169–196. Academic Press, 2008.
- D. Tasche. Estimating discriminatory power and PD curves when the number of defaults is small. Working paper, Lloyds Banking Group, 2009.